



# Disaster Recovery Best Practices

~~~~~

# Disaster Recovery Best Practices



## Preface

As the field of disaster recovery evolves and new expertise comes into being, new technologies and methodologies will reshape recovery strategies. Disaster Recovery Best Practices is intended to be an evolving reference and a compilation of contributions from many state agencies.

Your input on this reference is welcome. To contribute material or for questions and assistance, please contact:

Dietra Thomas  
Business Continuity Coordinator  
Division of the State CIO  
4430 Broad River Road  
Columbia, SC 29210-4012  
(803) 896-0177  
[djt@cio.sc.gov](mailto:djt@cio.sc.gov)

# Disaster Recovery Best Practices

~~~~~

## Table of Contents

	Page
Introduction.....	1
Disaster Recovery Plans .....	2
BEST PRACTICES .....	2
Risk Assessment (RA).....	3
BEST PRACTICES .....	3
Business Impact Analysis (BIA) .....	4
Recovery Class / Tier .....	5
BEST PRACTICES .....	5
Recovery Time Objective (RTO).....	6
BEST PRACTICES .....	6
Recovery Point Objective (RPO).....	7
BEST PRACTICES .....	7
Typical RTO's and RPO's.....	9
Sample Procedure to Build a Disaster Recovery Plan.....	10
Acknowledgements .....	14

# Disaster Recovery Best Practices

## Introduction

“Disaster Recovery” addresses that portion of a Business Continuity Plan which deals with the recovery of IT processing capabilities.

The measures and procedures put in place to provide disaster recovery are specific to:

- E The criticality of each processing system and its assigned RTO (Recovery Time Objective).
- E The tolerable data loss potential for that system and its assigned RPO (Recovery Point Objective).
- E The hardware, software, networking and other operating environment characteristics of the system and its dependencies.
- E The monetary and man-power resources of agencies.
- E The recovery management preferences of agencies.

These factors constitute valid and varied differences between disaster recovery strategies and can result in wide discrepancies in the disaster recovery environments designed by different agencies. However, since common components share disaster recovery considerations, agencies are urged to share their techniques and experience.

**Agencies are strongly urged to work together on their disaster recovery strategies, to pursue sharing backup and recovery facilities and resources, and to earnestly consider cooperative ventures.**

# Disaster Recovery Best Practices

## Disaster Recovery Plans

### BEST PRACTICES

- ✓ All IT facilities need documented Disaster Recovery plans.
- ✓ Copies of Disaster Recovery plans must be kept offsite and accessible to the recovery team.
- ✓ Disaster Recovery plans should be tested no less than once a year.
- ✓ Disaster Recovery plans must be maintained and should be reviewed for changes no less than once a year.
- ✓ Critical IT infrastructure requires Incident Response Plans (IRP), a type of Disaster Recovery plan specific to an infrastructure component, which specifies how to handle and recover from possible impacts that would impair that component's ability to deliver the necessary performance.
- ✓ Disaster Recovery plans must be supported by plans for all logistical support departments; such planning is contained in a Business Continuity Plan (BCP).
- ✓ Platforms which support distributed processing for one or more systems which require recovery should ideally plan for recovery at the same site. If different sites are chosen, then those sites should be sufficiently proximal to ensure the minimum throughput for each recovered system.
- ✓ If one or more related or co-dependent (front-end, back-end, etc.) IT facilities choose a given recovery site, then the other facilities sharing the co-dependency should consider choosing the same recovery site; co-dependent IT facilities should work jointly in developing their recovery strategies. Proximity not only reduces networking costs and transfer times but also reduces exposure to network disruption (fewer potential points) and recovery times.

# Disaster Recovery Best Practices



## Risk Assessment (RA)

A Risk Assessment identifies the threats to the business from natural or human-mediated (intentional or accidental) sources, rates the probability of the threats occurring, determines what impact each threat could have in consideration of the precautions and protections mounted against it, and produces a risk exposure factor for each threat, usually expressed as a probability of impact such as 1 in 10 chance of total loss, 5% probability of a power outage longer than 3 days, etc.

Risk exposures are primarily used to evaluate the degree to which a business should implement protection measures and how much investment is justified, especially for protecting structures which facilitate business processes such as buildings, power houses, network cables, communication towers, or other enabling facilities.

Risk Assessment data can be used in conjunction with the Business Impact Analysis to apply probabilities to business process outages. However, in terms of disaster recovery planning which deals with restoring IT business processes on which today's business environment is so highly dependent, it is generally accepted that IT has a zero risk exposure tolerance and so recovery planning is always required and the investment evaluation is based on the BIA alone (see "Business Impact Analysis").

### BEST PRACTICES

- ✓ A RA should be conducted for all IT enabling facilities such as data center buildings, power houses, and external communications facilities (network cables, relay stations, towers, etc.)
- ✓ Based on its RA, appropriate protection and impact mitigation measures should be implemented for each IT enabling facility.
- ✓ RAs should be reviewed for changes no less than once a year.

# Disaster Recovery Best Practices

## Business Impact Analysis (BIA)

A Business Impact Analysis (BIA) identifies each processing system's criticality, i.e. how much impact would outage of the system cause, and how long after the outage occurs would the impact be incurred. Criticality is used to plan the recovery to acceptable recovery requirements, and to determine how much should be spent on recovery capabilities, considering the following caveats or Rules Of Thumb (ROT):

**Rules Of Thumb:** (1) Technology can decrease recovery times & data loss exposure.  
(2) The faster the recovery, the more costly the technology.

Impact can be both quantitative or qualitative. Quantitative impact is usually expressed in dollars, e.g. loss of income, fines, loss of business base, etc. Qualitative impact is usually expressed as a non-numeric description, e.g. loss of lives, disruption of emergency services, damage to business reputation, loss of trained employees, missed business opportunities, etc. Impact, whether quantitative or qualitative, must be correlated with how long after the outage the impact will be incurred. Some impacts occur once at a specified time after the outage and others have recurring, and sometimes varying, impacts at various times after the outage.

The combined impact / time lapse determines the criticality of the system as illustrated in the following chart showing **Gartner's Sample Classification**:

Recovery Class/Tier	Financial Impact	Legal or Contractual	Service Impact	System Name
Multisite application	\$500,000 / day	No	Within 45 minutes	Order, Web
1	\$200,000 / day	No	Within 24 hours	Order, Internal
1	\$300,000 / day after 2 days	No	1 to 3 days	ERP
2	< \$100,000	Yes	5 to 10 days	Finance Reporting
3	None	No	Not time-critical	Data Warehouse

# Disaster Recovery Best Practices

## Business Impact Analysis (BIA) continued

### Recovery Class / Tier

Each recovery class (or tier) ranks the criticality of the system. The following basic criticality structure provides three criticality classes:

- 1 = HIGH:** the system must be recovered within a **short time**  
or **significant harm or cost** will be incurred.
- 2 = MEDIUM:** the system should be recovered within a **moderate time frame**  
or **some damage** will be incurred.
- 3 = LOW:** **little or no damage** will be incurred for an **extended period of time**.

### BEST PRACTICES

- ✓ All IT facilities need to conduct a BIA for all systems.
- ✓ BIAs are used to guide decisions on outage tolerance and how much to invest in reducing outage exposure.
- ✓ Based on these decisions, each system is assigned a Recovery Time Objective (RTO) and a Recovery Point Objective (RPO).
- ✓ BIAs should be reviewed for changes no less than once a year.

# Disaster Recovery Best Practices

## Recovery Time Objective (RTO)

Recovery Time Objective (RTO) is the target lapse of time after a disaster by when the system should be recovered. In other words, RTO is the maximum amount of time which can elapse between the point in time when a disaster destroys the service and the point in time by which the service must be recovered or unacceptable consequences will ensue.

RTO sets a target limit on recovery time and hence is used to guide decisions in planning how recovery from a disaster will be achieved. Recovery options are limited by how much expenditure is justifiable to achieve recovery in a given time. Generally, the faster the recovery, the more expensive the solution.

Recovery investment is a business decision determined by weighing the costs of lengthening outage periods (see BIA) against the increasing expenditures needed to shorten the outage period. A realistic RTO is one which can be met by methods which fall within the recovery investment limit.

### BEST PRACTICES

- ✓ RTOs are best indicated by a Business Impact Analysis.
- ✓ A realistic RTO is one that is achievable within expenditure limits.
- ✓ All systems should be assigned a RTO, even those with low criticality.

# Disaster Recovery Best Practices

## Recovery Point Objective (RPO)

Recovery Point Objective (RPO) is the target point of recovered work. This is the state of work which will be restored to the recovered system after a disaster. Work can only be restored to the point at which it was last saved and removed to safe-keeping before the disaster.

Potential data loss is calculated by adding the times between backups and the time lapse until the backup is stowed in a safe place. It is the sum of the time since the last backup was taken and when it is safely stowed offsite.

Consider, for example, the following scenario of a “weekly” backup:

- ↓ backup1 is taken on Friday, March 5
- ↓ backup1 tapes are packaged on Monday, March 8
- ↓ backup1 boxes are stowed in the offsite vault on Tuesday, March 9
- ↓ backup2 is taken Friday night, March 12
- ↓ backup2 tapes are packaged on Monday, March 15
- ↓ disaster destroys the data center at 03:05 am Tuesday, March 16
- ↓ the only backup available for restore is backup1, taken March 5 = **11 days previous**.

If the potential data loss is more than the desired RPO, then backup and storage procedures and timing should be adjusted accordingly. In general, the lower the RPO, the more expensive the solution to achieve it. Financial considerations can increase the tolerance for a higher RPO.

### BEST PRACTICES

- ✓ The frequency of backup creation is guided by the Recovery Point Objective (RPO).
- ✓ The procedure to stow backups offsite is guided by the RPO.

# Disaster Recovery Best Practices

## Recovery Point Objective (RPO) continued

### BEST PRACTICES continued

- ✓ Backups should be stored offsite in a location which is:
  - ↓ Suitable for the physical protection of the media and its contents.
  - ↓ Secure.
  - ↓ Accessible by disaster recovery teams.
- ✓ To improve the probability of a readable copy, keep at least two full backups in offsite storage, in addition to the full backup being taken and shipped to offsite storage.
- ✓ To ensure data integrity, a media retention plan should be developed and formalized where tape media is tracked during its life cycle. Retention and re-use rates should be based on the media's reliability metrics including length of life and number of uses. The purpose of this plan is to insure that media is retired before data is lost.

# Disaster Recovery Best Practices

## Typical RTO's and RPO's

Gartner's suggested **Business Process Service Levels** :

Class	Business Process Services	Service Levels				
		Scheduled Hrs x Days	Availability		RTO	RPO
			%	Downtime		
1 * (RTE)	<ul style="list-style-type: none"> <li>Customer-/ Partner-Facing</li> <li>Functions Critical to Revenue Production</li> </ul>	24 x 7	99.9 %	< 45 mins. / month	2 hrs.	0 hrs.
2	<ul style="list-style-type: none"> <li>Less-Critical Revenue-Producing Functions</li> <li>Supply Chain</li> </ul>	24 x 6 ¾	99.5 %	< 3.5 hrs. / month	8 – 24 hrs.	4 hrs.
3	<ul style="list-style-type: none"> <li>Enterprise Back-Office Functions</li> </ul>	18 x 7	99 %	< 5.5 hrs. / month	3 days	1 day
4	<ul style="list-style-type: none"> <li>Departmental Functions</li> </ul>	24 x 6 ½	98 %	< 13.5 hrs. / month	5 days	1 day

\* Class 1 application services are those with a RTE (Real-Time Enterprise) strategy and are those that the enterprise would suffer irreparable harm from if they were unavailable.

# Disaster Recovery Best Practices

## Sample Procedure to Build a Disaster Recovery Plan

1. Perform a Risk Assessment (RA) to identify the risk exposures. See “Risk Assessment (RA)” for more information on RA.
2. Use the results of the RA to determine and implement requisite protection and precaution measures.
3. Identify every application and the IT resources required to support it.
4. Perform a Business Impact Analysis (BIA) to determine the quantitative and qualitative cost per unit of time of application outage for all the applications. See “Business Impact Analysis (BIA)” for more information on BIA.
5. Determine how much expenditure can be justified to mitigate the outage costs identified in the BIA.
6. Determine the Recovery Time Objective (RTO) for each application. See “Recovery Time Objective (RTO)” for more information on RTO.
7. Determine the Recovery Point Objective (RPO) for each application. See “Recovery Point Objective (RPO)” for more information on RPO.
8. Design (or review) a backup methodology for the application to ensure the RPO can be met. Storage vendors and storage services can present available options which include:
  - a. Performing tape backups and transporting the tapes to an offsite vault.
  - b. Managing your own offsite storage facility (vault) or contracting with a storage service provider.
  - c. Performing backups directly to offsite tape.
  - d. Using dasd mirrors to enable taking tape backups with no (or less) application downtime.
  - e. Creating synchronous or asynchronous copies on offsite dasd.

# Disaster Recovery Best Practices

## Sample Procedure to Build a Disaster Recovery Plan continued

9. Design a recovery strategy for the application to ensure the RTO can be met. Recovery site providers and recovery service providers can present available options which include:
  - a. **Hot Site** – a fully serviced facility providing the necessary environment (A/C, power, water, cabling facilities, etc.) provisioned with all required hardware which is loaded, configured and ready to go.
  - b. **Warm Site** – same as a hot site but the software (OS/applications, etc) will need to be loaded and configured.
  - c. **Cold Site** - a fully serviced facility providing the necessary environment (A/C, power, water, cabling facilities, etc.) which will need to be provisioned with the required hardware.
  - d. **Mobile Site** – an IT facility which is delivered to a pre-determined recovery site and may, or may not, house the required hardware upon delivery.
  - e. **Hot Drop or Quick Ship** – an arrangement with a provider to deliver a hardware component within a pre-arranged time much shorter than normal; these arrangements provide for priority to be given to these orders upon short notice and typically contain provisions to shorten or circumvent delays associated with the usual procurement process.
10. Document the Disaster Recovery Plan ensuring that (1) the plan will be accessible after a disaster, and (2) procedures are put in place to maintain the plan.

Note: The plan will be more current and useable (and its maintenance easier and less frequent) if titles, positions or functions are used in the main body of the plan while citing specific names only in appendices and where the documentation is person-specific, such as contact lists.

The plan documentation should include:

- a. Specific recovery procedures sufficiently detailed that they could be implemented by someone with the appropriate skill set but no knowledge of the agency or its functioning.

# Disaster Recovery Best Practices

## Sample Procedure to Build a Disaster Recovery Plan continued

10. b. An action plan detailing who is responsible for what and when, including:
  - ↓ who assesses the situation and what criteria are used,
  - ↓ who declares disaster and the procedures involved,
  - ↓ who builds the recovery environment and the procedures involved,
  - ↓ who comprises the recovery teams and who are the alternates,
  - ↓ who activates the recovery teams and the notification procedures,
  - ↓ who manages funding and other procurement needs,
  - ↓ who manages the recovery process, resolves problems and conflicts, and makes management decisions, and
  - ↓ what the reporting structure is, complete with contact numbers.
- b. All support documentation including:
  - ↓ Contracts and other legal documents.
  - ↓ Graphical summaries (maps, charts, diagrams, etc.).
  - ↓ Technical references, guides, procedures and other documentation.
  - ↓ Reference information such as directories, inventories, indices, and other 'look-up' references.
  - ↓ Pre-printed forms or other process defining tools.
  - ↓ Contact information for
    - (1) recovery team members and recovery managers,
    - (2) employees and their emergency contacts (next of kin),
    - (3) normal providers, alternate providers, and providers of recovery services,
    - (4) hardware servicing and software support,
    - (5) customers and users,
    - (6) local, county, state and federal emergency services,
    - (7) governing bodies, related agencies and other stake holders.

# Disaster Recovery Best Practices

## Sample Procedure to Build a Disaster Recovery Plan continued

11. Design and implement procedures to test the plan and apply updates. It is desirable to have different people man the tests so that as many people as possible are familiar with the details of the plan and the recovery process; this improves the likelihood of having experience available for an actual recovery.
12. Design a method for detecting and applying changes to keep the plan current. This is critical to ensuring that the plan will be effective when it is needed; constant change is a business reality, for example, consider how frequently a business must update its telephone list.
13. The entire plan should be exercised no less than once a year; portions may be exercised independently more frequently, especially to verify modifications. This process checks for changes, verifies if expectations are still realistic, and provides the opportunity to train employees and reinforce plan knowledge.
14. Monitor business changes that could impact the plan. Organizational changes may impact departmental interfaces or affect the way logistical support is provided. A location on which the plan depends on may no longer provide the expected facility. Provider agreements may change procurement plans. It is important to remain mindful of the plan dependencies and watch for any changes affecting those dependencies which could adversely impact the plan.

# Disaster Recovery Best Practices

## Acknowledgements

AOC subcommittee, Systems Management Services, January – December 2004

Doug Mader, DOR (Chair) – End-User Support

Russell Collier, DHEC – Network Events Monitoring

Debbie Dunn, CIO – Asset Management

Lee Foster, DOT – Software Distribution

Dave Gerth, CIO – Problem/Change/Configuration Management

Eric Pinckney, CIO (CIO support group liaison)

Floyd Stayner, DNR – Network Events Monitoring

Dietra Thomas, CIO – Disaster Recovery

John Ward, DOC (AOC Coordinator)

Charlie Zeberlein, CIO - Performance and Capacity Management